REPORT NO. 8004
AUGUST 1980

Donald A. Norman

ERRORS IN HUMAN PERFORMANCE

CENTER FOR HUMAN INFORMATION PROCESSING
LA JOLLA, CALIFORNIA 92093

UNIVERSITY OF CALIFORNIA, SAN DIEGO

This research was supported by the Advanced Research Projects Agency under Contract N00014-79-C-0515, NR 157-434, under terms of ARPA Order No. 3767, and was monitored by the Office of Naval Research, Personnel and Training Research Programs. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency, of the Office of Naval Research, or the United States Government.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: 8004	2. GOVT ACCESSION NO.: AD-A091 925	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle): Errors in Human Performance		5. TYPE OF REPORT & PERIOD COVERED: Final Report June 1, 1979 – May 31, 1980
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s): Donald A. Norman		8. CONTRACT OR GRANT NUMBER(s): N00014-79-C-0515
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Center for Human Information Processing, University of California, San Diego, La Jolla, California 92093		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: NR 157-434
11. CONTROLLING OFFICE NAME AND ADDRESS: Personnel and Training Research Programs, Office of Naval Research (Code 458), Arlington, Virginia 22217		12. REPORT DATE: August 15, 1980
		13. NUMBER OF PAGES: 44
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report): UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Human information processing
Human error
Human-machine interface
System design
Attentional memory
Short-term memory

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Work on this contract examined various applications of human information processing, focusing on the analysis of human error. As a result of the contract investigations, several areas of potential application were identified. One area of immediate possible application is in the design of human-machine interfaces, with particular concern for aspects of system design that lead to human error, taking into account what is known about human attentional and short-term memory limitations. The major focus of the research effort was devoted to the study of human error, especially those made by highly skilled operators in complex, high-demand systems. A classification scheme for these errors was developed.

Errors in Human Performance

Donald A. Norman

Program In Cognitive Science

Center For Human Information Processing

University of California, San Diego

This research was sponsored by the Advanced Research Projects Agency and the Personnel and Training Research Programs, Psychological Sciences Division, Office of Naval Research.

Contemporary views of the information processing structure of the human have developed with considerable precision. Although the full picture of human processes is far from being understood, sufficient knowledge does now exist to approximate many human functions. Thus, even though the appropriate theoretical structure is still under considerable study and debate, engineering approximations can be developed that capture actual performance of people in specific tasks with sufficient accuracy for use in applications.

It is well known that short-term memory poses severe limitations on a person's ability to perform complex tasks. And, although the theoretical picture and status of STM is still under debate, several engineering statements can be made with reasonable certainty:

1. The capacity of STM is between 5 and 10 items.

2. Information within STM can be searched as an approximate rate of 100 msec/item.

3. Separate STM functions apply for different forms of encoding: verbal, motor, pictorial, spatial.

4. Information can be retained in STM through rehearsal, but rehearsal itself is interfered with by other activities.

These simple statements apply to situations in which a person must use STM in order to accomplish a task. Examples of such situations are the entering of numerical settings into equipment after being told by radio (voice) what values to use. Examples of these activities include pilots and air traffic controllers who must set such things as atmospheric pressure, altitude settings, sourse headings, and radio frequencies,

oftentimes while doing other tasks.

In a similar way, it appears that certain properties of the human attentional system should be susceptible to approximations that state the capacity limits of the attentional system, the sets of activities that interfere with one another (that draw from the same attentional resources) and the sets of tasks that do not interfere. Unfortunately, attentional limitation suffers from the lack of a unit in which to make the statements, so that no quantitative assessment of performance can yet be made.

Work on this contract was intended to be a one year, exploratory effort. This is a pilot approach to the development of an applied discipline. In the limited work time available, focus was directed toward the study of performance errors by humans in a variety of situations.

Studies of Human Error

Accidents often occur without errors, and errors often occur without accident. Even when error leads to accident, it is usually the case that the accident was multi-determined, that numerous human and environmental incidents combined to cause the accident. As a result, real accidents are difficult to categorize. Errors, however, are more tractable. In the sections that follow, we discuss a possible theoretical and empirical approach to the study of error. It is important to realize that these studies are in their infancy. Despite the importance of accidents (and the importance of human error in the causation of accidents), there is surprisingly little work on human error, almost

none at the theoretical level. One authority on human factors and engineering psychology (John Senders, personal communication) has just conducted an extensive bibliographic search and has found essentially no literature. The extensive analyses of nuclear power plant errors collected by Swain at Sandia Associates (personal communication) provides an empirical base of errors, but the analysis is entirely empirically oriented, with a deliberate attempt to avoid theoretical interpretations. (As a result, the categorization alone -- without the errors -- is several pages long and is descriptive rather than predictive. Moreover, it is primarily useful for the situation for which it was intended and does not readily generalize to other situations.) The analyses presented here are only suggestive of the techniques that will be followed.

One side aspect of this phase of the research is that it coincides with major new developments in human-machine interfaces, with computer and CRT-based display systems just starting to be introduced. Our initial investigations of errors in a computer environment indicate that certain errors are more likely to occur here than in other environments (e.g., "mode errors"). This work should contribute, therefore, to the design and applications of these new control systems.

The term "human error" includes a number of different kinds of incidents. It is useful to distinguish among the sources of human error. First, we need to say something about the genesis of human performance and the stages of processing that are involved. Then, we can identify the sources of different kinds of error. For this classifica-

tory purpose, a simplified view of performance is sufficient. An important dividing line beween two major classes of error is the formation of an intention to take some action. Errors up to and including the formation of the intention are called mistakes. Errors in the performance of the intention are called slips. That is, the person is in some situation that has to be recognized (through the use of perceptual, problem-solving and decision-making processes). Then, given the situational analysis, the person must determine what action is to be taken (through the matching of the current situation with previous experience, coupled with decision-making and problem-solving processes). We call the highest level of specification of that action the "intention." Now, once the intention is formed, it controls a hierachical assemblage of action schemata that eventually lead to the control of human output devices (limbs, voice control, eye movements), and a physical response is made.

Types of Errors:

Mistakes -- Errors in the formation of an intention

Slips -- Errors in the execution of an intention

Errors That Result from a Lack of Knowledge

The inexperienced or incompletely trained person. Lack of knowledge can be of two types. First, the person may be inexperienced or incompletely trained. This can lead to both mistakes and slips: mistakes when through insufficient knowledge the inappropriate intention was

formed; slips when through insufficient knowledge the actions were not performed properly. Neither of these cases is of particular interest, for the cause and the remedy is clear: better training. (Just how to provide proper training is not so clear, but this is clearly a separate topic.)

The well trained, skilled person. The second form of error that results from a lack of knowledge occurs when the person is well trained and skilled, but where full knowledge of the situation is not available, either because of faulty system design or because of problems internal to the person (such as mental overload). These errors are almost entirely in the classificaion "Mistakes." These mistakes can potentially be avoided.

Errors Even Though There Is Full Knowledge

Mistakes. Even when the person has full information of the state of the situation, mistakes and slips can occur. Mistakes arise when the situation is misclassified, or when inappropriate decisions and response selections are made. One major source of such errors has been amply categorized by workers in the decision making literature (in particular, Kahneman and Tversky, who have shown the emphasis on "typicality" and "representativeness" and those who are from what we will call the "Oregon School" of decision theorists who have shown the human inability to combine data in appropriate ways). A second source of mistakes is from errors in the retrieval and use of memory information (leading to what we have called "description mistakes," a name taken from our analysis of memory retrieval problems, Norman & Bobrow, 1979). One other major form

of mistakes is "system induced error" which we review shortly.

Slips. Finally, even if the correct intention is specified and the person skilled and well practiced, slips in performance can arise. Here there are numerous possible places for error. Our analysis of slips indicates that they can be classified into a number of types: check list errors, mergers, misorderings, captures, and intrusions (Norman, 1980). We suspect these slips follow the counterintuitive rule: the more skilled the person, the more likely the slip.

System Induced Errors

It is not sufficient to analyze error by an analysis only of the information processing stages within the person. The person is working within a system, a system which has task demands, environmental demands, social and societal demands. These demands can often be overriding in their determinants of the action that is to be performed. Moreover, people themselves do not operate in isolation. The human is part of the system, with numerous social and emotional aspects. Part of the environment is a large number of human artifacts -- our technology -- much of which serves as important adjuncts to our processing capabilities, sometimes purposely, as when we use calculators, hand written notes, check lists, or charts and instruments, and sometimes incidentally, as when we place objects in a pile in a specific location as reminders of the tasks we are to do (Norman, 1979).

Response compatability. Some forms of system induced errors come about from the design of the equipment, wherein certain responses are

"natural," others not so natural. The difficulty occurs when it is the unnatural response that may be required. This particular issue of design has received considerable attention in the human factors and ergonomics literature. Suprisingly, however, although notions such as "response compatability" are well understood at the level of practice in terms of the general factors that contribute to a set of displays and instruments being "compatible" or not, the underlying processing mechanisms that cause one system to be compatible while another is not are not well understood. Moreover, it is well known that training has dramatic effects on compatibility, so that situations that once appeared to be incompatible can be perceived after sufficient experience to be quite natural and comfortable. It is clear that the development of an appropriate "internal model" can change the compatibility structures, and although this is talked about in the literature, to our knowledge the theoretical underpinnings have never been explored. There is a cri-ical need for such basic level understanding at this time, for conventional system design is undergoing major changes with the substitution of computer displays for conventional instruments, computer control for manual control (with the human becoming a supervisor rather that an operator), and with the introduction of advanced visual, auditory, and manipulatable displays and controls.

Supervising rather than operating. A major change is coming about in the control of systems: operators are now more like management than like skilled manipulators of controls. Pilots are not so much piloting as managing, ships are controlled by sophisticated computer systems, nuclear plants are monitored by computer: the role of the human

operator is changing. In emergency situations, the demands upon the human are high, and the environment conducive to errors of all sorts. For one, the large complex systems of today are not well designed for large failures. Alarms are not part of a system design, but tend to proliferate with each individual need for a warning. In times of danger, the large number of alarms that are active can add to the already high mental work load. System models are not readily available. The operator must be inventive in a situation that often makes inventiveness difficult. Too little is known about the use of mental models by operators, too little about how operators must divide their efforts. What information must be provided the operator? In what form? Much of what is presented today appears to be that which is really required. Does a pilot really need to know outlet temperature at the jet nozzle? Probably not: the pilot is really attempting to assess the state of the engine, and some more global, integrated measure might very well be more appropriate.

Errors resulting from social interaction. The Tenerife air crash resulted largely from deficits in social interactions: between crew members, between the crew personnel and the Air Traffic Control personnel (coupled with time pressures on the KLM crew). The relatively large number of incidents labelled "Controlled flight into terrain" (Seigal, 19--), in which commercial aircraft crash while the crew is so preoccupied with fixing a minor problem that no one flies the aircraft, is an example of a problem in social interaction. The Eastern Airlines crash in Miami is an example (Dec. 29, 1972). The problem was the failure of the landing gear light to go on, and the attempt to determine the cause

so preoccupied the entire crew that they neglected to monitor their altitude. These and similar incidents result in part from the insistence of the pilot on being in complete charge of all aspects of the situation, whereas in fact, he ought to delegate all aspects and take an overall, supervisory role. Too many people attempting to help get in the way.

Similar incidents occur when too many people must divide up responsibility for handling an incident. Unless the specific operations required of each participant are spelled out in great detail, there is apt to be ambiguity over the division: the right engine catches on fire; the co-pilot turns off the alarm, the pilot shuts down the engine. No one pulls the fire extinguisher handle (in this case, each thought the other had done it). Note that it is probably impossible to spell out divided responsibilities in detail: with complex situations, one cannot predict all the possible modes of failure, let alone all possible combinations of equipment states.

Another form of social situation occurs when there are higher level pressures and demands on the participants in a situation, often of the form that are never really voiced explicitly, sometimes not quite even known to the people involved. It is "unmanly" to admit to being afraid, "unmanly" to admit error. In Southern California, amateur scuba divers have died, sometimes in quite shallow water, sometimes after struggling in the water. Yet, often these divers fail to release the weight belt (which can have as much as 20 pounds of lead), sometimes even fail to release their heavy catch of fish, lobster, and abalone. A problem in

that when divers do release their weight belts and, therefore, survive, their colleagues are apt to laugh (for once the diver is safe, who is to know whether it was really necessary to release the weight belt or not).

Military pilots are reported to have delayed exiting from their disabled and burning aircraft while they reset switches or memorize panel readings in order to prepare for the accident hearings. The plant operators at Three Mile Island were more concerned about preventing damage to the equipment than about safety, and at one point turned off the high-pressure injection emergency pumps (that had been triggered automatically by the plant situation), an act they had been told to do in other situations to avoid damage to the equipment.

In Golder's analysis of P-3 pilot errors, he found that social factors did not alwas correspond to safety factors. Thus,

> TAXI OFF THE RUNWAY AT NIGHT and MAKE A WHEELS UP PASS are not as likely to be tolerated by cognizant Navy officials as EXCEED DESIGN AIRSPEED FOR FLAP SETTING and FAILURE TO NOTE VOR OFF FLAG. ... the error perceived to be the number one career wrecker, TAXI OFF THE RUNWAY AT NIGHT, compared with the third from the least (11th in rank order) perceived career wrecker, SALVAGE A LANDING FROM A POOR APPROACH THAT SHOULD HAVE BEEN A WAVE-OFF, tells us several interesting things. If a pilot runs his aircraft off the taxiway at night, he feels it will affect his career, it will rattle him somewhat, it will embarrass him ... On the other hand, to continue on a landing ... that should have been abandoned ... is number one in the fun ranking (Golder, 1978).

As Golder puts it, "The task then is to change pilots' perceptions of the 'system' and their attitudes, in order to make it more career damaging and less fun to commit the errors that are likely to lead to loss of lives and aircraft damage."

<u>Summary of the Existing Literature on Error</u>

The existing literature on error divides causes and sources of error into a number of different classifications. Thus, Kidd (1962) suggested errors occur as failure to detect a signal, incorrectly identifying a signal, as incorrect value weighting, as errors in action section, and as errors of commission. De Greene (1970) has a similar analysis. Welford (1973) suggested that errors occur when the human operator reaches the limit of its capacities or when it has received inadequate information. This led Welford to propose a four-fold categorization of errors: ignorance, speed, span of apprehension, and the presence of random activity. Singleton (1973) recognized that social factors are often involved.

Meister and Rabideau (1965) suggest a category of out-of-sequence performance, failure, incorrect, and non-required performance. Singleton (1976) suggested that errors are either perceptual (the operator's mental model is inadequate or wrong) or motor (timing mismatches and sequencing disorders).

As the previous section indicates, our own analyses of a reasonably large number of human errors do not lead to the same forms of categorization. For one thing, with modern, complex systems, it is not really possible to categorize errors into unique classes. Seldom does there appear to be a single cause for an error. Rather, there appears to be a combination of social, situational, and cognitive factors that interact, often with numerous erroneous actions, eventually the total causing sufficient difficulties that an accident occurs.

Tenerife

An excellent example of the multiple causes of accidents is the collision of the Pan American 747 and the KLM 747 at Tenerife, March 27, 1977. (The following analysis comes from the ALPA report: Roitsch, Babcock, & Edmunds, 1979.) A number of different factors contributed to the crash, no single one being sufficient to have triggered the accident.

1. Both aircraft crews had been on duty for a long time period.

2. The KLM crew was concerned about duty time, and was worried about not being able to return to Amsterdam without changing crews and putting passengers up in (insufficient) hotel space.

3. The weather was closing in fast.

4. The Pan Am flight was ready to go an hour before KLM, but had to wait because it couldn't clear the taxi-way until the KLM plane moved out of the way.

5. The pilot of the KLM flight was the chief pilot of KLM, with strong opinions about flying, but who had in actuality few duty hours as an operational pilot (he was mostly involved in training). The KLM co-pilot had been recently checked out for the 747, by the pilot.

6. The communication with Air Traffic Control (ATC) was not optimum and there is evidence that the Pan American flight gave up trying to change its runway assignment because of this problem.

7. There was confusion as to the point at which the Pan Am aircraft should leave the runway (to a taxi-strip, thereby permitting the KLM plane to take off). The ATC said the third exit, but this was not possible (the required turn was too sharp), and so Pan Am, after several attempts at clarification, evidently assumed it was the fourth exit that was meant.

8. The KLM pilot attempted to take off without tower clearance, but was stopped by the co-pilot. The KLM plane then told the tower that it was " ... now ready for takeoff and we are waiting for our ATC clearance." The tower responded with the ATC clearance, and the KLM plane acknowledged the clearance and took off. However, the tower acknowledgement was not for takeoff, only for the flight plans.

9. The tower did not stop the takeoff, but rather asked Pan Am to state when it was clear of the runway.

10. Fog prevented the KLM plane and the Pan Am plane from seeing each other, or the tower from seeing either plane.

These factors all intermixed to cause the incident. No single one was responsible.

Clearly, a model of human performance which will allow us to understand the source of accidents such as these must exceed the traditional bounds of human information processing models. We must include in our models all of the kinds of factors we have listed above. Accidents often arise out of the configuration of a large number of these kinds of factors. Clearly, no simple model could account for our observations. We must see how these kinds of systems interact -- how they interact to lead to errors and how they interact to lead to appropriate performance most of the time. We wish to build simulation models consisting of a set of interacting systems of the sort outlined above. Such models will allow us to observe the effects of errors in one subsystem on the error rate in the overall system. The details of such models are not yet clear to us. We believe, however, that the time is ripe to begin other similarly operating models. Moreover, we are convinced that such models are essential if we are to understand the complex phenomena involved in

most real world accident situations (the 'interactive' aspect of modeling). We have already begun to explore in the context of comprehension and perception (see the section on perception for an example). We hope to be able to develop models which will allow us to talk about interactions among larger units and among individuals -- each driven by their own goals and each deciding on their own actions to best satisfy those goals. The multi-goal models discussed in the context of goal achieving systems might well form a prototype for such a model. Obviously, we have to gather a good deal more information about these kinds of situations and a lot more experience with methods of modeling such situations before we can develop anything with the precision we would like -- nevertheless, this is the direction we feel we must go if we are to be successful.

Three-Mile Island

Although the common interpretation of the Three-Mile Island incident places a large part of the blame on "human error," careful analysis of this accident makes this seem not so simple. (The following is based upon the analysis presented in the special issue of the IEEE Spectrum, 1979.)

Essentially, the accident was triggered by a blocked line, apparently the result of transferring resins from the demineralizer in the secondary coolant system, a common operation. A resin block developed causing water backup which tripped the condensate pumps. This would cause the main feedwater pumps to trip, which in turn would stop the flow of steam through the turbine, causing the automatic sensors to

trip the turbine. All of this happened in about one second. So far, this is a normal accident and no danger exists.

Because the turbine shut down, the auxiliary feedwater pumps started up, and reactor power was cut back (by pushing the control rods into the core). However, two valves were closed in the auxiliary feedwater system, thus preventing auxiliary feedwater from replacing the main feedwater. These valves were shut a few days earlier during servicing and they were not turned back on (this was probably a sequencing error, triggered by a side effect condition). Moreover, due to sloppy panel design and the use of large tags to signal out-of-service equipment, the fact that these valves were closed was not easily determined by the operators.

The electromatic relief valve that had opened automatically to relieve pressure in the reactor coolant system should have closed (13 seconds into the incident). Indeed, it was instructed to close, and its indicator on the control panel indicated that it was closed. However, the valve was stuck open. (The indicator monitored the control signal rather than the actual state of the valve, another design error.)

From here on, there are a lot of different actions and analyses. Our main point, however, is that although there was technically sufficient information for the operators to determine that the relief valve was still open, this would have required considerable debugging on their part, checking out a mental model that was implausible. The operators thought the valve was shut, the indicator said it was shut, and one sign of disagreement (high temperature in the valve leading to the drain

tank) was consistent with an alternate model: that there was a slight (previously known) leak in the relief valve. A critical indicator that possibly would have triggered re-assessment of the model of the reactor state was the drain tank pressure indicator. However, this was located behind the 7 ft. high primary control room panel, the panel which housed most of the critical instruments. (Panel layout in these control rooms is inexcusably bad: see the Lockheed EPRI report.)

In all the analyses that we have seen of this incident, no one has considered the mental models that must be constructed by operarors of complex systems to determine the true state of the system. The instruments are, at best, secondary measurers of system state. Sometimes, the instruments are themselves reading derived signals. The mental work load is high.

This work can benefit directly from the work from our previous contract research and the related work of other ONR-sponsored groups studying the use of mental models. This work was part of a program concentrating on the training of skills, but it is immediately and directly relevant here. The system is faulty, not the operators. The mental model demanded of the operators is complex. Moreover, the immediate evidence is consistent with one state of the model, whereas the actual systm is at another state. Were the initial evidence inconsistent, then the task would have been more easily solved. It seems clear that a different set of system monitors that are more in tune with the mental models used by the system operators. Borrowing from the experience of the ONR research contracts on Interactive Instructional (tutorial) sys-

tems (which includes our own work) it should be possible to develop systems that can check for consistency of model states, flagging the operator that something is inconsistent.

The Analysis of Human Error

The preceding arguments suggest that error associated with complex systems is apt to be a multi-faceted phenomenon, one not readily analyzable as a single, simple categorizable set of system: the person, the social interaction, the task demands (which may include subtle social pressures). At the least, the analysis must include a careful consideration of the three systems involved: cognitive, physical, and social.

A Sample Error Analysis: Enabling Analysis

Consider the following analysis of a commonplace task: using the Xerox 4500 copier/sorter system, equipped with a Rusco "copier control device" (that requires a plastic card and the entering of a budget number on a ten key keyboard in order to energize the copier). This machine is simple enough both to reveal the complexities of the analysis and to cause human error of several sorts. We have performed controlled observations of operators of this machine. In this section, we review briefly the forms of errors that have been observed, describe an *enabling* analysis of the machine operation that reveals some of the potential for error, and then discuss how the analysis procedure needs to be extended and applied to a wider variety of tasks.

The required sequence of operations is reasonably complex, and the following forms of errors have been observed.

Enable machine

Enter plastic card for budget number

ERROR: fail to enter card

ERROR: enter card incorrectly (several times)

Enter budget number on keyboard, check number and enter confirmation

ERROR: enter number wrong

ERROR: 10% of sample ignores opportunity to check budget number and enters confirmation without checking

Turn machine ON

Set proper machine state (sorter, 1 or 2 sided copying, light original, auxiliary paper tray, choice of paper sizes and types)

ERRORS: machine is checked primarily when it is desired to do nonstandard operation. Else, assumption is that machine is set properly in default settings. Often a false assumption. Some unnecssary checking does occur for error conditions (paper in auxiliary tray) that would be signaled if checking were required.

Set proper number of copies

ERROR: if only 1 copy desired, fail to check.
if more than one copy desired, there is high likelihood of checking (as above: check mostly for non-standard operation)

Place original on platen

ERROR: fail

ERROR: place original on previous original

ERROR: misorient original

push "start" button or

push "document assist" button

ERROR: push wrong button for task

This set of operations now cycles for number of copies desired. There are a number of special conditions here and numerous potentials (and realizations) of errors. Rather than prolong the analysis, however, consider only the cleanup phase.

cleanup machine state

remove copies from sorter or paper tray

ERROR: remove only some copies (usually due to interruption during performance of task)

remove original from platen

ERROR: fail

remove originals from document-assist tray

ERROR: fail

remove plastic budget card from slot

Frequent ERROR: leave card in slot

take all materials out of copier room

ERROR: take only subset of materials

Categorization of Common Copier Errors

The most common errors that we have observed in the copying operation can be classified as these:

default error

(assume default conditions apply, without checking)

mode error

(believe task to be in mode i when actually in mode j)

place keeping error

(lose track of place in sequence-- sometime leading to sequence-list error, sometimes to repetition error)

sequence-list error

(fail to enter budget card, to remove original, to remove card)

repetition error

(repeat step in sequence: put second original on platen)

cleanup error

failure to clean up from side-effects. Not an error in terms of task performance, but often a serious cause of accident.

(Entering budget card into slot has side-effect that card is in slot. To correct, must remove card when task is finished.)

description error

perform operation similar to that desired, but erroneous

(use wrong control, put wrong original on platen)

These categories are rough, for this work is just beginning. Moreover, as we have already stated, it is false to expect errors to be neatly

categorizable: errors of complex systems have multiple causes. Nonetheless, even this simple categorization has obvious correlates in other situations.

Human Error in Operational Situations

Consider the following errors, observed in operational situations:

Failure to re-open valves closed during test (Three-Mile Island incident). Senders (1979) reports that the observed frequency of failure to re-open vlaves is sometimes as high as 0.01. This is a cleanup error. Closing the valve was necessary to do the test, with the side effect that the valve is now closed, but no longer on critical (enabling) path for the task of testing. Failure to cleanup is, in part, a sequence-list error, but it is technically not an error for the task at hand, only for a subsequent task (ie., emergency operation of the reactor).

Landing with landing gear up: Sequence list error

Pressurizing refueling system while probe is extended: mode error

Land at wrong airport (or at wrong runway, or on taxi strip rather than runway). These probably have multiple causes, based around description errors. Landing at the wrong airport occurs through perceptual confusions (among other things), where airport descriptions match (personal communication: Private pilot who landed at Miramar Naval Air Station). Similar reasons for landing on taxi strip (Palomar Airport, CA).

Use wrong control: often description error. Can also be mode error or activation error (see Norman, 1980).

Act upon expectations rather than actual situation: default error, plus well known effects of expectations on cognitions. Look for lowered landing gear and report seeing them, even though they were

not down (personal communication, Navy pilot); assume take-off clearance has been granted, when it was not (Tenerife KLM-Pan Am incident).

Analytical Techniques

To analyze a given task, we need to develop some analytical machinery. Several potential tools are available, and we propose to extend these analytical methods. Promising approaches include:

1. Enabling analysis. A technique which we are developing to specify the set of enabling conditions for any operation. [An example follows.]

2. Petri net analysis. This is a form of occurrence analysis that is used in system evaluation (see Holt, 1971; Peterson, 1977). This method has both diagrammatic appeal and formal analytic power. Although we have only begun our assessment of this technique, it appears to be a superset of enabling analysis. [An example follows.]

3. Coupling. When multiple systems run quasi-autonomously (the human, the system that is to be controlled, the environment), operations in one system may not have a formal coupling with operations in another, and so these interactions will not necessarily appear on an enabling or Petri net analysis. Often, the coupling is assumed to be the duty of the human operator, but unless we have a method of making the coupling explicit, both the analysis is weakened and, in the case of real systems, there is a potential for accidents to occur. [Example: The following P-3 errors (from Golder, 1978) are examples of a lack of coupling in that the required operation had no immediate coupling to the system in terms of enabling further operation. Rather, the coupling was mental, in that the pilot was expected to know that the operation was essential for safe operation (or for operations far removed in time from the required action): failure to remove pilot covers before takeoff, takeoff with flaps not set at "takeoff and approach," restart an engine in-flight with circuit breakers not properly set. In similar fashion, landing with the gear up, or failure to turn off an automobile light when leaving the automobile are examples of coupling situations where there is no immediate coupling of the act and the system performance.]

4. Side effects. Operations or actions may have results that do not affect the performance of the desired task, but that may be

deleterious for future tasks. These outcomes are side effects, and an analysis of their occurrence important.

Enabling Analysis: An Example

Consider the operation of the Xerox machine. In particular, consider just one small segment of the operation, the use of a plastic card to establish permission to use the machine and to identify the budget to be charged for accounting purposes. The principle underlying the basic enabling analysis is the specification for each goal that is to be accomplished, those conditions that enable the goal, those that are required to reach the goal, those that inhibit (prevent) the goal from coming about, and the side effects that result from the operations. Our work on the analysis of situations is still at an early stage of development, and so the analyses to be presented here are designed primarily to show the potential for these analyses not the fully worked out details.

Here is a simplified analysis of the task of making a single side copy of an original document on a Xerox machine. The analysis is shown in Figure 1 and in the following statements.

Goal: have copy of one page document

Requires: make copy of document

Requires: remove copy from machine

Requires: machine cycle

Goal: make copy of document

Requires: original on platen

Requires: machine cycle

Goal: original on platen

Requires: place original on platen

Requires: have original

Requires: platen free

Requires: physical constraints be met

SIDE EFFECT: original is out of sight, on platen

Goal: machine cycle

Requires: machine operation

Requires: push START button

Enabled by: ready light

Requires: machine on

Inhibited by: current machine cycle

error in machine state

Enabled by: entering of budget number

Goal: enter budget number

Requires: enter 4-digit budget number

Requires: knowing (remembering) budget number

Enabled by: enter card into slot

Requires: possession of card

SIDE EFFECT: Card is out of sight, in slot

Inhibited by: card in wrong orientation

illegal card

Requires: enter number confirmation

Enabled by: "enter 4-digit budget number"

COUPLING: visual confirmation of number

The analysis shown here (and in Figures 1 and 2) is meant only to be suggestive of the direction which our work is moving. This analysis still has some difficulties. There is not yet a clear understanding about the relationship of an enabling condition to a required one. "Inhibited by" conditions have similar difficulties. The distinctions among machine operations and states, and human operations, states, knowledge, and interactions are not well done. Moreoever, this analysis is much simplified. It is interesting to note what it does <u>not</u> require. For one, the default settings of the machine were assumed to be correct (although often they are not). For another, the goal to get an acceptable copy does not require anywhere in it a requirement to remove the original from the machine: a side effect and coupling problem. Couplings are not shown well by this analysis. For that purpose, let us turn to a Petri net analysis, shown in Figure 2.

Arcs enter and leave transitions and places. The rule is that a transition "fires" whenever all its inputs are alive. Consider that each place that is active is marked with a token. If all the inputs to a transition are marked, then the transtition fires. When this happens, all the tokens responsible for the firing are removed and tokens are put in the places indicated by the outputs of the transition.

Petri nets have some powerful analytical properties, mostly revolving around the concept of "reachability:" a given configuration of markings reachable by the network. This allows for determination of

such things as whether a system will thrash, or block (deadlock). For the current analysis, the Petri net forces a complete determination of the human-machine-environmental-interaction. It illuminates couplings, those areas where there is no physical requirement on the system, but rather a mental requirement that is supposed to be satisfied before operation can continue. Not surprisingly, this is where errors occur. Thus, the Xerox machine operates if the start button is pressed and the ready light is on. There does not need to be an original on the platen. The budget number unit requires that a four digit budget number be entered, but there is no requirement that it be a legitimate budget number. The budget card, the original, and the copies have nothing dependent upon their removal from the machine, so there is apt to be failure to take all originals, or all copies, or the budget card.

Petri nets are also useful to examine timing relations to see whether there are critical race conditions. There are none in this particular example.

Required Work on Analytical Techniques

The formal methods of analysis described so far -- Petri nets and enabling analysis -- have both virtues and deficits. Petri nets do not do well at determining enabling conditions. Neither technique seems good at determining couplings, the effect of mental work load, or side effects. It seems clear that considerable work is needed to determine appropriate analytical techniques. We feel we have made a start in this direction. Two other techniques are also promising, one that we have helped develop, the other from work in artificial intelligence.

Problem solving and planning spaces are another possible source of relevant information. Thus, the work of Sacerdoti (1975) and Schmidt, Sridharan & Goodson (1978) on planning spaces, and the work at SRI on robot planning and deduction (summerized in Nilsson, 1980) offer good potentials. Interestingly, though, these techniques are riddled with side-effects and couplings, for they are designed for computer implementations in which every critical condition is checked before an operation is performed and perfect memory an computation is assumed. The checking condition is not consistent with our observations of human performance, and perfect memory and computation is certainly not true for humans. Thus, the work in Artificial Intelligence is more limited than we had hoped.

In our work on human error (ONR Technical Report 7903) we analyzed an extensive collection of naturalistic errors, categorizing them according to a theoretical analysis. This analysis emphasized "Slips," where the intention was correct but the action did not carry out the desired intention. That classification is presented here as Table 1.

Table 1

A Classification of Slips Based upon Their Presumed Sources

I. Slips that result from errors in the formation of the inention.

A. Errors that are not classified as slips: errors in the determination of goals, in decision making and problem solving, and other related aspects of the determination of an intention.

B. Mode errors: erroneous classification of the situation.

C. Description errors: ambiguous or incomplete specification of the intention.

II. Slips that result from faulty activation of schemas.

A. Unintentional activation: when schemas not part of a current action sequence become activated for extraneous reasons, then become triggered and lead to slips.

1. Capture errors: when a sequence being performed is similar to another more frequent or better learned sequence, the latter may capture control.

2. Data-driven activation: external events cause activation of schemas.

3. Associative activation: currently active schemas activate others with which they are associated.

B. Loss of activation: when schemas that have been activated lose activation, thereby losing effectiveness to control behavior. This leads to such slips as:

1. Forgetting an intention (but continuing with the action sequence).

2. Misordering the components of an action sequence.

3. Skipping steps in an action sequence.

4. Repeating steps in an action sequence.

III. Slips that result from faulty triggering of active schemas.

A. False triggering: a properly activated schema is triggered at an inappropriate time, leading to:

1. Spoonerisms: reversal of event components.

2. Blends: combinations of components from tow competing schemas.

3. Thoughts leading to actions: triggering of schemas meant only to be thought, not to govern action.

4. Premature triggering.

B. Failure to trigger; when an active schema never gets invoked, because:

1. The action was preempted by competing schemas.

2. There was insufficint activation, either as a result of forgetting or because the initial level was too low.

3. There was a failure of the trigger condition to match, either because the triggering conditions were badly specified or the match between occurring conditions and the required conditions was never sufficiently close.

<u>Mode errors</u>. Mode errors occur when a system behaves differently depending on the state it is in, and the user's action is inappropriate for the current state. A typical situation which gives rise to mode errors is the use of a computer text editor. Here, the distincition between "command" mode and "input" mode is not usually well marked, resulting in responses appropriate to one mode being entered while in the other. (Another example of a likely occasion for mode errors occurs in the interpretation of the indicators on a Heads-Up Display (HUD) where at times the same display can mean different things depending upon the mode of the aircraft.) It is clear that mode errors are more likely where there is a task that has different interpretations of the same responses, depending upon the mode, and where the modes are not easily distinguishable from one another. There are four major factors which determine the likelihood of the mode error:

1. varying the similarities of the states;

2. marking the states explicitly;

3. changing the goodness of the cueing function that marks the state to serve as a good cue for the appropriate response for that state; and

4. varying the similarity of the responses required within each mode. When completely different response sets are required, the distinctiveness between the states is heightened, making mode errors quite unlikely.

Description errors. Description errors result from ambiguous or incomplete specification of the intention. There are two essential classes of description errors, one arising from memory retrieval, the other arising from action specification. Description errors will occur in situations with context-dependent descriptions, with time and processing pressures and with multiple responses being required, some of which have identical descriptions when viewed out of context.

Capture errors. Capture errors occur when a sequence being performed is similar to another more frequent or better learned sequence, and the latter captures control.

Data-driven errors. A data-driven error occurs when an external event leads to initiation of an action. A classic example of data-driven errors occurs in the Stroop phenomenon.

Associative activation errors. These errors occur when currently active schemas activate others with which they are associated. In many ways, these are similar to data-driven errors, and so the inducing situation is closely related, except that in this case the intruding stimuli need not be of the same form as the information required for the task. Rather, the intruding stimuli must be highly associated with information that is of the same form required by the task. These errors occur with reasonable frequency during typing.

Forgetting the intention (but continuing with the action sequence). In this situation, the action sequence continues apparently normally, but the reason for the action has been forgotten. The situation is

revealed when the action sequence is completed and the person discovers that he or she has no idea what should be done next. A common situation that gives rise to a loss of intention error is where once the desired goal and requisite action sequence is determined, some other action sequence must first be performed in order to get ready to do the desired sequence: call this the "preparatory sequence." If interference occurs during the preparatory sequence, there is apt to be forgetting of the goal state.

Skipping a step in an action sequence. Leaving out a step in an action sequence is most often caused by a memory failure, often by a combination of distraction and heavy memory load. To avoid this, the situation must be designed so that the exact position in the action sequence can be deduced by examination of the current state. The size of the action component that is forgotten varies. Thus, if the action sequence is hierarchically structured, then the amount of action sequence that is lost depends upon the exact point in the hierarchy where the forgetting takes place.

Basically, a major cause of step skipping, we suspect, is that a long action sequence is interrupted, and then, in the resumption of that sequence there are insufficient clues to determine the exact state of the completion of the sequence, or at least not without considerable effort. A typical situation would occur in the following of a checklist for the setting up of a panel for appropriate configuration for a desired action. If the setup is interrupted, then the place on the check list cannot easily be determined by examination of the panel: it

must be remembered.

Repeating steps in an action sequence. This set of errors actually derives from the same sources as the preceeding class, skipping steps in an action sequence. Essentially, it results from losing one's place in the sequence.

Slips that result from faulty triggering: Spoonerisms, reversal of event components. Experimental generation of Spoonerisms and related errors has been performed by Baars and Motley (Baars & Motley, 1976; Baars, in press). These slips are elicited by generating sequence conflicts, often with prior activations that generate competing action plans, and often with time and processing pressures.

Blends. A blend results when several actions are in conflict, no final decision about which to perform has yet been made, and time pressures demand immediate action. In this case there is a tendency for the resulting action to be a blend or combination of the competing actions.

Thoughts leading to actions. In this situation, mental thoughts interfere with ongoing actions, often leading to the performance of something that was meant only to be thought, not to be done. The experimental situation requires that a person be required to do two tasks, one overt, the other mental. Thus, if someone were required to name the objects in a complex pictorial display while simultaneously keeping track of the number, we would expect that with sufficient time and processing pressures, the numerical count would occasionally intrude upon the primary task.

Premature triggering. In this situation highly salient or important actions occur in advance of their desired time. The likelihood of premature triggering probably increases the more difficult the sequence and the higher the time and processing pressures.

Failure to trigger. One frequent cause of failure to trigger an appropriate action schema is the absence of an appropriate triggering condition in the environment. Thus, in our observations of people using a Xerox copier, some obligatory actions are skipped (thereby leading to failure of the next step). The skipping appears to result from the lack of observable cues or "forcing functions" that would trigger the action. (Examples: failure to place the plastic card for the accounting charges into the appropriate slot, or failure to remove the original from the platen at the completion of the task.)

Suggestions Towards the Development of Design Principles

Our studies of errors, skills, human performance, and perception point the way towards the development of a set of design principles. Complex systems need to be designed with the considerations of the human operator as a fundamental target of the design. This tends not to be true today, in part because designers are not presented with appropriate design principles that they can use during the design stage. Rather, human factors are usually incorporated afterwards, when a design is reviewed by a set of human engineers and human factor experts. This is too late. We propose work towards development of a set of design principles. The goal is to give designers tools that can be used during the

design phase itself. Although it is premature to state these tools, some of the basic principles are sufficiently well-developed from our research efforts that they can be stated here. The concepts of a mental model, cueing and blocking functions, and intention-based systems, which are referred to in these principles, are discussed in more detail in the later part of this section. The principles are as follows:

1. Establish a mental model to be used by the user, and design the system around this mental model. Spell out the mental model in detail, being explicit about the assumptions. Design all displays and operations to be directly consistent with this model, minimizing the transformation required between the actual system and the user's internal mental model.

2. Observe human processing limits. Minimize short-term memory load. Minimize attentional distraction and attentional overload. But, keep the operator continually up-to-date as to the status of all states of the internal model. This means the operator must continually be observing and interacting with the system in a meaningful way.

3. Do an analysis of the cognitive load on the operator, including demands on short-term memory and on attentional resources.

4. Design around error points. Provide cueing and blocking functions where the side effects and coupling demands require them.

5. Use intention-based systems. Make the system understand the user. Make the system responsive to the needs and capabilities of the user.

These principles are not as yet well worked out. They do give some hint as to the direction in which we propose the research to go towards the specification of design principles based upon fundamental principles of performance and perception.

Cueing and blocking functions. Systems should be designed so that the set of responses potentially available for a particular situation is as limited and constrained as possible.

A blocking function is defined to be an event that "blocks" a response. (This is the function of "interlocks" in good human factors design.) Thus, the fact that a copying machine may not work unless the paper has been loaded or the correctbudget number been entered into the appropriate device is a blocking function. However, the lack of action by the machine poses few constraints upon the set of alternative corrective action. Thus, a blocking function prohibits continued operation until a desired action sequence has been accomplished, but it is up to an appropriate cueing function to indicate to the operator exactly which action it is that should be performed.

Studies of mental models. When a human engages in action, the choice of the action, the details of the exact specification of the control sequence, and the outcome all result from the person's interaction with the environment, the response of the intermediary system with which the person is interacting, and the details and timing of the act itself. In the selection and guidance of an action, a person must have an internal model of this combined system, although this internal model is often made up of a set of smaller models which may be partially inconsistent.

Considerable work on the development and use of mental models occurred in several projects examining intelligent computer assisted instructional systems, where models of the student were important components of the instructional system, including the works of Burton &

Brown (1979), Miller (1979), Goldstein (1979), and Stevens and Collins (1977). Similar approaches were used in our studies of intelligent computer assisted instructional systems (see Gentner, 1979). In the study of "response compatability" it has been observed that different compatible relations can be formed when subjects use different internal mappings of response to action (see any standard human factors book, e.g., McCormick, 1976). We believe this to be an important observation and the abilities of people to develop models that make certain mappings of response to action natural, perhaps making other mappings unnatural, should be explored.

One interesting comparison is between the model that the users have of the system with which they are interacting and the model the system has of its users. In general, we find that systems shortchange the users, failing to recognize the the particular powers and needs of the human operators but instead requiring of them information in ways that are most useful to the system itself. As systems become more complex, the requirement that humans conform to the machine structure becomes more and more unrealistic: What happens is that machines require that humans act like machines rather than doing the necessary translation allowing people to provide what is convenient and natural for the people. Forcing people to interact on the machine's terms is not only inconvenient -- more importantly, because it is an unnatural mode of interaction, it is a primary cause of human error.

References

Anderson, J. R. Language, memory, and thought. Hillsdale, N.J.: Erlbaum Associates, 1976.

Baars, B. J. On eliciting predictable speech errors in the laboratory: Methods and results. In V. Fromkin (Ed.), Errors of linguistic performance: Slips of the tongue, ear, pen, and hands. New York: Academic Press, 1980.

Baars, B. J., & Motley, M. T. Spoonerisms as sequencer conflicts: Evidence from artificially induced errors. American Journal of Psychology, 1976, 83(3), 467-484.

Bartlett, F. C. Remembering. Cambridge: Cambridge University Press, 1932.

Bellman, K. The conflict behavior of the lizard, sceloporus occidentales, and its implication for the organization of motor behavior. Unpublished doctoral dissertation, University of California, San Diego, 1979.

Bobrow, D. G., Kaplan, R. M., Kay, M., Norman, D. A., Thompson, H., & Winograd, T. GUS, a frame driven dialog system. Artificial Intelligence. 1977, 8, 155-173.

Broadbent, D. E. Word-frequency effect and response bias. Psychological Review, 1967, 74, 1-15.

Burton, R. R., & Brown, J. S. An investigation o computer coaching for informal learning activities. International Journal of Man-Machine Studies, 1979, 11, 5-24.

Collins, A. M., & Loftus, E. F. A spreading activation theory of semantic processing. Psychological Review, 1975, 82, 407-428.

deGreene, K. B. (Ed.). Systems psychology. New York: McGraw-Hill, 1970.

Gentner, D. R. Toward an intelligent automated tutor. In H. F. O'Neil, Jr. (Ed.), Procedures for instructional systems development. New York: Academic Press, 1979.

Golder, T. V. P-3 Pilot Errors: A conceptual approach. Center for advanced research, Naval War College, June, 1978.

Goldstein, I. P. The genetic graph: A representation for the evolution of procedural knowledge. International Journal of Man-Machine Studies, 1979, 11, 51-77.

Hewitt, C., Bishop, P., & Steiger, R. A universal modulator ACTOR formalism for Artificial Intelligence. Paper presented at the Third International Joint Conference on Artificial Intelligence, Stanford, August, 1973.

Hirst, X. X., Spelke, X. X., Reaves, X. X., Caharack, X. X., & Neisser, X., X. Dividing attention without alternation or automaticity. Journal of Experimental Psychology: General, 1980, 109, 98-117.

Hollerbach, J. M. A study of human motor control through analysis and synthesis of handwriting. Unpublished doctoral dissertation, Massachusetts Institute of Technology, 1978.

Holt, A. W. Introduction to occurrence systems. In E. L. Jacks (Ed.), Associative information techniques. New York: American Elsevier, 1971.

Kidd, J. S. Human tasks and equipment design. In R. M. Gagne (Ed.), Psychological principles of system development. New York: Holt, Rinehart & Winston, 1962.

Kucera, H., & Francis, W. N. Computational analysis of present-day American English. Providence, R.I.: Brown University Press, 1967.

LaBerge, D., & Samuels, S. J. Toward a theory of automatic information processing in reading. Cognitive Psychology, 1974, 6, 293-323.

Leeuwenberg, E. L. J. A perceptual coding language for visual and auditory patterns. American Journal of Psychology, 1971, 84, 307-349.

Luce, R. D. Individual choice behavior. New York: Wiley, 1959.

McClelland, J. L. On the time relations of mental processes: An examination of systems of processes in cascade. Psychological Review, 1979, 86, 287-330.

McClelland, J. L., & Rumelhart, D. E. An interactive activation model of the effect of context in perception, part I. ONR technical report (in press) and CHIP Technical Report, Center for Human Information Processing, University of California, San Diego. La Jolla, California, 1980.

McCormick, E. J. Human factors in engineering and design (4th ed.). New York: McGraw-Hill Book Company, 1976.

Meister, D., & Rabideau, G. F. Human factors evaluation in system development. New York: Wiley, 1965.

Miller, M. L. A structured planning and debugging environment for elementary programming. International Journal of Man-Machine Studies, 1979, 11, 79-95.

Newell, A. Production systems: Models of control structures. In W. C. Chase (Ed.), Visual information processing. New York: Academic Press, 1973.

Nilsson, N. J. Principles of artificial intelligence. Palo Alto, California: Tioga Publishing Co., 1980.

Norman, D. A. The role of active memory processes in perception and cognition. Proceedings of the XXI International Congress of Psychology, Paris, 1976. To be published.

Norman, D. A., & Bobrow, D. G. On data-limited and resource-limited processes. Cognitive Psychology, 1975. 7, 44-64.

Norman, D. A., & Bobrow, D. G. On the role of active memory processes in perception and cognition. In C. N. Cofer (Ed.), The structure of human memory. San Francisco: Freeman, 1976.

Norman, D. A. Slips of the mind and an outline for a theory of action. (Tech. Rep. No. 88). La Jolla, California: Center for Human Information Processing Technical Report, November, 1979.

Norman, D. A. Errors in human performance (Tech. Rep. 8004). University of California, San Diego, July 1980.

Norman, D. A. Post Freudian slips. Psychology Today, April, 1980.

Norman, D. A. Human error. Reader's Digest, 1980, in press (various editions: American, French, ...) (This is a simplified version of the Psychology Today article.)

Norman, D. A. Categorization of action slips. Psychological Review, in press.

Norman, D. A., & Shallice, T. Attention to action: Willed automatic control of behavior. 1980 (in preparation).

Palmer, S. E. Fundamental aspects of cognitive representation: Theory and metatheory. In E. Rosch (Ed.), ?. Hillsdale, N. J.: Erlbaum Associates, in press.

Peterson, J. L. Petri nets. Computing surveys, 1977, 9, 223-252.

Posner, M. I., & Snyder, C. R. R. Attention and cognitive control. In R. L. Solso (Ed.), Information processing and cognition: The Loyola symposium. Hillsdale, N. J.: Erlbaum Associates, 1975.

Restle, F. Theory of serial pattern learning: Structural trees. Psychological Review, 1979, 77, 481-495.

Rosenbloom, P. S. The APS Reference Manual. Unpublished manuscriipt, University of California, San Diego, 1979, (also Carnegie-Mellon University).

Rubenstein, E. (Ed.). Special issue on Three Mile Island and the future of nuclear power. IEEE Spectrum, 1979, 16, (11), whole issue.

Ruffel-Smith, H. P. A simulator study of the interaction of pilot work-load with errors, vigilance and decisions (NASA Technical Memorandum 78482). Washington, D.C.; NASA Scientific and Technical Information Office, 1979.

Rumelhart, D. E. A multicomponent theory of the perception of briefly exposed visual displays. Journal of Mathematical Psychology, 1970, 7, 191-218.

Rumelhart, D. E. A multicomponent theory of confusion among briefly exposed alphabetic characters (Tech. Rep. 22). San Diego, Ca.: University of California, San Diego, Center for Human Information Processing, 1971.

Rumelhart, D. E. Notes on a schema for stories. In D. G. Bobrow and A. M. Collins (Eds.), Representation and understanding: Studies in cognitive science. New York: Academic Press, 1975.

Rumelhart, D. E. Understanding and summarizing brief stories. In D. La Berge and S. J. Samuels (Eds.), Basic processes in reading: Perception and comprehension. Hillsdale, N. J.: Erlbaum Associates, in press.

Rumelhart, D. E. Introduction to human information processing. New York: John Wiley & Sons, 1977.

Rumelhart, D. E., & Norman, D. A. Accretion, tuning and restructuring: Three modes of learning. In R. Klatsky and J. W. Cotton (Eds.), Semantic factors in cognition. Hillsdale, N. J.: Erlbaum Associates, in press.

Rumelhart, D. E., & Ortony, A. The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), Schooling and the acquisition of knowledge. Hillsdale, N. J.: Erlbaum Associates, 1977.

Rumelhart, D. E. & Siple, P. The process of recognizing tachistoscopically presented words. Psychological Review, 1974, 81, 99-118.

Sacerdoti, E. A. The non-linear nature of plans. Paper presented at the Fourth International Joint Conference on Artificial Intelligence, Tblisi, USSR, 1975.

Saltzman, E. Levels of sensorimotor representation. *Journal of Mathematical Psychology*, 1979, *20*, 91-163.

Schmidt, R. A. A schema theory of discrete motor skill learning. *Psychological Review*, 1975, *82*, 225-260.

Schmidt, R. A. The schema as a solution to some persistent problems in motor learning theory. In G. E. Stelmach (Ed.), *Motor control: Issues and trends*. New York: Academic Press, 1976.

Schmidt, C. F., Sridharan, N. S., & Goodson, J. L. The plan recognition system. *Artificial Intelligence*, 1978, *11*, 45-83.

Shaffer, L. H. Latency mechanisms in transcription. In S. Kornblum (Ed.), *Attention and performance IV*. New York: Academic Press, 1973.

Shiffrin, R. M., & Schneider, W. Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 1977, *84*, 127-190.

Simon, H. A., & Kotovsky, K. Human acquisition of concepts for sequential patterns. *Psychological Review*, 1963, *70*, 534-546.

Singleton, W. T. Theoretical approaches to human error. *Ergonomics*, 1973, *16*, 727.

Singleton, W. T. Skill and accidents. In *Occupational accident research*. Stockholm: Gotab, 1976.

Welford, A. T. *Causes of human error*. Paper presented to Royal Australian College of Surgeons, 1973.